

Decomposing a Graph Into Expanding Subgraphs

Guy Moshkovitz*

Asaf Shapira†

Abstract

A paradigm that was successfully applied in the study of both pure and algorithmic problems in graph theory can be colloquially summarized as stating that *any graph is close to being the disjoint union of expanders*. Our goal in this paper is to show that in several of the instantiations of the above approach, the quantitative bounds that were obtained are essentially best possible. Three examples of our results are the following:

- A classical result of Lipton, Rose and Tarjan from 1979 states that if \mathcal{F} is a hereditary family of graphs and every graph in \mathcal{F} has a vertex separator of size $n/(\log n)^{1+o(1)}$, then every graph in \mathcal{F} has $O(n)$ edges. We construct a hereditary family of graphs with vertex separators of size $n/(\log n)^{1-o(1)}$ such that not all graphs in the family have $O(n)$ edges.
- Motivated by the Unique Games Conjecture, Trevisan and Arora-Barak-Steurer showed that given a graph G , one can remove only 1% of G 's edges and thus obtain a graph in which each connected component has good expansion properties. We show that in both of these decomposition results, the expansion properties they guarantee are (essentially) best possible even when one is allowed to remove 99% of G 's edges. In particular, our results imply that the eigenspace enumeration approach of Arora-Barak-Steurer cannot give (even quasi-) polynomial time algorithms for unique games.
- Sudakov and the second author have recently shown that every graph with average degree d contains an n -vertex subgraph with average degree at least $(1 - o(1))d$ and vertex expansion $1/\log^{1+o(1)} n$. We show that one cannot guarantee a better vertex-expansion even if allowing the average degree to be $O(1)$.

The above results are obtained as corollaries of a new family of graphs, which we construct by picking random subgraphs of the hypercube, and analyze using (simple) arguments from the theory of metric embedding.

1 Introduction

In recent years a certain paradigm has emerged which roughly says that any graph is close to being a vertex-disjoint union of expanders. Unlike the ϵ -regular partitions of Szemerédi [30] and the cut decompositions of Frieze-Kannan [9] which are relevant only for dense graphs (i.e., with $\Theta(n^2)$ edges), this paradigm is applicable for graphs of arbitrary density. Generally speaking, its

*School of Mathematics, Tel Aviv University, Tel Aviv, Israel 69978. Email: guymosko@tau.ac.il. Supported in part by ISF grant 224/11.

†School of Mathematics, Tel Aviv University, Tel Aviv, Israel 69978. Email: asafico@tau.ac.il. Supported in part by ISF Grant 224/11 and a Marie-Curie CIG Grant 303320.

usefulness stems from the fact that it allows for a divide-and-conquer approach, in essence reducing a problem on general graphs to the special case of expander graphs.

For algorithmic problems, this paradigm has been useful, for example, in designing approximation algorithms related to the Unique Games Conjecture [14], as in the seminal work of Arora-Barak-Steurer [3] and of Trevisan [31]. It has also seen applications in property-testing algorithms [7, 10], as well as in data structure design [25]. In graph theory, instantiations of this paradigm include the theorem of Lipton-Tarjan-Rose [22] on vertex separators in hereditary families, the results of Linial-Saks [21] and Leighton-Rao [19] on low-diameter decompositions, as well as results in the field of graph minors [17, 27] and cycle packing [6]. See Section 2 for some further discussion.

We note that in all applications of the regularity lemmas [9, 30] one uses the same notion of expansion¹. On the other hand, when dealing with sparse graphs and applying the above-mentioned paradigm, each application calls for a different notion of expansion. Our main goal in this paper is to show that in several of the above-mentioned applications, the different notions of expansion that were used are quantitatively best possible.

As it turns out, all our results can be deduced from a construction of a new family of graphs (stated in Theorem 1), whose main property is that every subgraph of it has a small *edge separator*, which in particular implies small vertex/edge expansion. We next recall some basic definitions.

Definition 1.1 (Edge/Vertex Separator). *An edge separator in a graph $G = (V, E)$ is a set of edges whose removal disconnects G into two (not necessarily connected) subgraphs, each on at most $\frac{2}{3}|V|$ vertices. We write $\text{sep}_E(G)$ for the minimum cardinality of an edge separator in G . A vertex separator is defined in the same way, only now vertices are removed rather than edges. We write $\text{sep}_V(G)$ for the minimum cardinality of a vertex separator in G .*

Note that $\text{sep}_V(G) \leq \text{sep}_E(G)$, simply by taking an arbitrary endpoint of each separator edge. The *girth* of a graph is the minimum length of a cycle in it. We henceforth write $\log(\cdot)$ for $\log_2(\cdot)$. Our main technical result in this paper, stated next, implies that there exist graphs with a super-linear number of edges, whose girth is $(\log n)^{1-o(1)}$ and whose t -vertex subgraphs all have expansion at most $(\log t)^{-1+o(1)}$.

Theorem 1. *For any n, k with $2 \leq k \leq \frac{1}{648} \log \log n$ there is an n -vertex graph $G = G_{n,k}$ satisfying:*

- i. G has average degree at least k and maximum degree at most $6k$.*
- ii. G has girth at least $\log n / (6k)^2$.*
- iii. Every t -vertex subgraph H of G with $t \geq \log n / (6k)^2$ satisfies*

$$\text{sep}_V(H) \leq \text{sep}_E(H) \leq \frac{t}{\log t} \cdot (\log \log t)^2. \quad (1)$$

Let us briefly explain why the quantitative estimates in Theorem 1 are (essentially) best possible. First, observe that for subgraphs H of $G_{n,k}$ in Theorem 1 on fewer than $\log n / (6k)^2$ vertices, we of course have a much stronger bound on $\text{sep}_V(H)$ and $\text{sep}_E(H)$ than the one stated in (1) for larger subgraphs. To see this, one just has to recall the following well-known fact.

¹In the setting of the regularity lemma, expansion is referred to as being ϵ -regular.

Fact 1.2. *Every tree T satisfies $\text{sep}_V(T) \leq 1$.*

So item (ii) in Theorem 1 implies that subgraphs H of $G_{n,k}$ on fewer than $\log n/(6k)^2$ vertices satisfy $\text{sep}_V(H) \leq 1$ and thus also $\text{sep}_E(H) \leq \Delta(H) \leq \Delta(G) (\leq \log \log n)$.² Note that the bound on $\text{sep}_E(H)$ cannot be improved by much beyond $\Delta(G)$ since G contains $K_{1,\Delta(G)}$ (the star with $\Delta(G)$ leaves) as a subgraph, and it is not hard to see that $\text{sep}_E(K_{1,\Delta}) \geq \Delta/3$. Second, as we discuss in Section 6, any family of graphs satisfying item (iii) (or variants of it) contains at most $n(\log \log n)^{O(1)}$ edges, so one cannot construct graphs as above with average degree much larger than $\log \log n$. Finally, it follows from Theorem 6 below that *any* graph of average degree at least 3 has a subgraph that is not a tree and does not contain a separator of size much smaller than the one in (1), so our estimate on the size of the separator in item (iii) cannot be reduced by more than some $\log \log t$ terms. Actually, the proof of Theorem 1 shows that in (1) we can replace the $(\log \log t)^2$ term by $(\log \log t)^{1+o(1)}$, but we chose to use the simpler/cleaner expression.

1.1 Small set expansion

Next we describe a strengthening of Theorem 1 which will be important for some of our applications. Notice that Theorem 1 gives a graph whose every subgraph has a *large* subset (i.e., consisting of at least $1/3$ of its vertices) that does not expand. One may thus further ask for a graph that does not contain even *small set expanders*; that is, a graph whose every subgraph has a small non-expanding set (i.e., consisting of $o(1)$ -fraction of the vertices). The notion of small set expanders has recently received much attention in theoretical computer science (see, e.g., [3, 28, 29]). For example, it is believed that solving unique games is intimately connected with the problem of finding small non-expanding sets.

For a graph H and a subset $\emptyset \neq A \subseteq V(H)$ we write $\partial_H(A)$ for the set of edges of H with precisely one endpoint in A , and $\phi_H(A)$ for the *edge expansion* of A in H , that is,

$$\phi_H(A) = \frac{|\partial_H(A)|}{|A|}.$$

We have the following “small set” counterpart of Theorem 1. Note that here we in fact prove the stronger property that one can essentially *partition* the subgraph into small non-expanding subsets.

Theorem 2. *Let $G = G_{n,k}$ be a graph from Theorem 1 with $2 \leq k \leq \frac{1}{648} \log \log \log n$. For every t -vertex subgraph H of G with $t \geq \log n/(6k)^2$, and every $1/t \leq \mu \leq 2/3$, there are $\lceil 1/4\mu \rceil$ mutually disjoint subsets $A_i \subseteq V(H)$ of size $\mu t/3 \leq |A_i| \leq \mu t$ satisfying*

$$\phi_H(A_i) \leq \frac{\log(1/\mu)}{\log t} \cdot (14 \log \log t)^2. \quad (2)$$

Note that in Theorem 2 we can find, in any large subgraph H , subsets that simultaneously have “measure” $o(1)$ and edge expansion $o(1)$. For example, by setting $\mu = 1/\log t$ in (2), each of the subsets $A \subseteq V(H)$ produced by the theorem satisfies $|A|/t = 1/\log t$ and $\phi_H(A) \leq (1/\log t)(14 \log \log t)^3 = 1/(\log t)^{1-o(1)}$.

² $\Delta(\cdot)$ denotes the maximum degree of a graph.

Similarly to Theorem 1, the quantitative estimates in Theorem 2 are (essentially) best possible. It follows from Theorem 6 below that any graph of average degree at least 3 has a t -vertex subgraph whose every subset of measure μ has edge expansion at least $\frac{\log(1/\mu)}{\log t}(\log \log t)^{-2}$. Thus, our estimate of the edge expansion cannot be reduced by more than some $\log \log t$ terms. Finally, note that while Theorem 2 only deals with subgraphs H on at least $\log n/(6k)^2$ vertices, any smaller subgraph is a tree (due to item (ii) in Theorem 1), implying that it trivially has small subsets with weak expansion.

1.2 Paper overview

In Section 2 we describe the main results of this paper, which are the applications of Theorems 1 and 2, showing that many decomposition-type results that were used in different areas of research are essentially best possible. These include applications in approximation algorithms, property testing, extremal problems in graph theory as well as the nested dissection method. The proof of Theorem 1 is given in Section 3 (with some technical proofs deferred to Section 5) and the proof of Theorem 2 is given in Section 4. Section 6 contains some concluding remarks and open problems. Let us thus give an overview of the proofs of our theorems.

The main idea of the proof of Theorem 1 is to start with the Boolean hypercube Q_d and choose an appropriate random subgraph G of it. Using (simple) probabilistic and combinatorial arguments, we show that this random subgraph G satisfies items (i) and (ii) of Theorem 1. To prove the third item, we first prove that all bounded degree subgraphs of Q_d have small edge expansion. To this end, we use an argument similar to the one used by Linial, London and Rabinovich [20] in the setting of metric embedding. Then, we proceed to “boost” this fact in order to show that every sufficiently large subgraph of G actually has a small edge separator, thus establishing item (iii) of Theorem 1.

In order to prove Theorem 2, we first show that the hereditary nature of item (iii) of Theorem 1 can be used to iteratively find smaller and smaller subsets while controlling their expansion. This enables us to find a single small subset having small expansion. Next, we remove this subset and look for another small subset, this time having small expansion in the remaining subgraph. We continue in this manner until we get almost a partition of the vertex set. Finally, we show that many of the subsets we obtained have small expansion also in the original subgraph.

2 Applications of Main Results

2.1 Edge density of hereditary families with small separators

A family of graphs is said to be *hereditary* if it is closed under taking induced subgraphs. The well-known Planar Separator theorem of Lipton-Tarjan [23] asserts that any n -vertex planar graph has a vertex separator of cardinality at most $O(\sqrt{n})$. This influential result led to many extensions for other hereditary families of graph (such as minor-free families, see [2]). The notion of vertex separator in graphs has since found numerous applications in studying both pure and algorithmic problems in graph theory (see [26] and its references). One of the first applications of separators was given by Lipton, Rose, and Tarjan [22] in their work on the nested dissection method. Intuitively,

their result states that the reason planar graphs (and more generally minor-free graphs) have linearly many edges is that they have small separators.

Theorem 3 ([22], Theorem 10). *For every $\epsilon > 0$ there is $C > 0$ for which the following holds. Let \mathcal{F} be a hereditary family of graphs such that every n -vertex graph $G \in \mathcal{F}$ satisfies $\text{sep}_V(G) \leq n/(\log n)^{1+\epsilon}$. Then every n -vertex graph in \mathcal{F} has at most Cn edges.*

We mention that Fox and Pach [8] strengthened Theorem 3 by proving that even separators of size $n/\log n(\log \log n)^{1+\epsilon}$ guarantee³ that every graph in \mathcal{F} has $O(n)$ edges.

Using Theorem 1 we next show that the separation requirement in Theorem 3 cannot be improved much beyond $n/\log n$.

Corollary 2.1. *There is a hereditary family of graphs \mathcal{F} such that every n -vertex graph $G \in \mathcal{F}$ satisfies $\text{sep}_V(G) \leq (n/\log n)(\log \log n)^2$, yet there is no $C > 0$ such that every n -vertex graph in \mathcal{F} has at most Cn edges.*

Proof. Let \mathcal{F} be the family of graphs defined as follows:

$$\mathcal{F} = \{G : G \text{ is an induced subgraph of } G_{n, \log \log n / 648} \text{ for some } n\},$$

where $G_{n, \log \log n / 648}$ are the graphs from Theorem 1. Note that \mathcal{F} is, by definition, a hereditary family of graphs. By items (ii) and (iii) of Theorem 1 (and Fact 1.2), every n -vertex graph $G \in \mathcal{F}$ satisfies $\text{sep}_V(G) \leq (n/\log n)(\log \log n)^2$, as required. On the other hand, by item (i) of Theorem 1, the family \mathcal{F} contains, for every large enough n , an n -vertex graph with $\Omega(n \log \log n)$ edges. This completes the proof. \square

2.2 Decomposing a graph into edge expanders

Trevisan [31], in his work on approximating constraints satisfaction problems related to the Unique Games Conjecture, gave a combinatorial algorithm for approximating unique games. The main component of the algorithm is a decomposition of any graph into disjoint expanders. Given the graph underlying the unique game input, the algorithm proceeds to solve the problem on each expander separately, which, as Trevisan shows, can be done in polynomial time. We note that we will revisit this paradigm for solving unique games in Section 2.3, where we discuss the well-known Arora-Barak-Steurer [3] algorithm. We begin with the formal definitions.

A (linear) *unique game* is a system of linear equations of the form $x_i - x_j = c_{i,j} \pmod{q}$. Let $\text{UG}_q(\delta)$ denote the computational problem of deciding, given a unique game with equations modulo q , whether there is an assignment satisfying at least a $(1 - \delta)$ -fraction of the equations or rather every assignment satisfies at most a δ -fraction of the equations (we are guaranteed that one of the two cases holds). The influential Unique Games Conjecture⁴ of Khot [14] postulates that $\text{UG}_q(\delta)$ is NP-hard for every $0 < \delta < 1/2$ and sufficiently large q . Trevisan's algorithm efficiently solves a weaker version of $\text{UG}_q(\delta)$ (and therefore does not refute the Unique Games Conjecture).

A graph $G = (V, E)$ is said to have *edge expansion* α if every subset $S \subseteq V$ with $|S| \leq |V|/2$ satisfies $|\partial(S)| \geq \alpha |S|$ (recall that the *edge boundary* $\partial(S)$ of S is the set of edges with precisely

³The result of [8] actually shows that having separators of size $n/\log n \log \log n(\log \log \log n)^2$ suffices.

⁴The version stated here was shown in [15] to be equivalent to Khot's original formulation in [14].

one endpoint in S). Trevisan proved the following decompositions theorem for arbitrary graphs. Roughly speaking, it asserts⁵ that, given any graph, one can remove few of the edges in order to disconnect it into disjoint graphs of high expansion⁶.

Theorem 4 ([31], Lemma 10). *Let $0 < \epsilon \leq 1/2$. From any n -vertex graph one can remove at most an ϵ -fraction of the edges to obtain a graph whose every connected component has edge expansion at least $\epsilon/12 \log n$.*

Using Theorem 1 we show in Corollary 2.3 below that, already for $\epsilon = 1/2$, the expansion guarantee in Theorem 4 cannot be improved much beyond $1/\log n$. In fact, we show that the same conclusion holds even when $\epsilon = 1 - o(1)$, that is, even if one is allowed to remove all but a $o(1)$ -fraction of the edges of the graph! To this end, we will first need the following observation regarding the graphs $G_{n,k}$ from Theorem 1.

Claim 2.2. *Any subgraph of $G_{n,k}$ with average degree at least 4 has at least $n^{1/72k^2}$ vertices.*

Proof. Let H be a t -vertex subgraph of G with average degree at least 4. Recall the well-known fact that the girth of a t -vertex graph of average degree at least 4 is at most $2 \log t$ (in fact, average degree at least 3 suffices, see [1]). Item (ii) in Theorem 1 thus implies $\log n / (6k)^2 \leq 2 \log t$, completing the proof. \square

Corollary 2.3. *Let $G = G_{n,k}$ be a graph from Theorem 1. Any subgraph of G with average degree at least 4 has a connected component with edge expansion at most $(1/\log n)(15k \log \log n)^2$.*

Proof. Consider a subgraph of G with average degree at least 4, and note that it has a connected component H with average degree at least 4. Writing t for the number of vertices of H , Claim 2.2 implies that $\log t \geq \log n / 72k^2$. It follows from item (iii) in Theorem 1 that the edge expansion of H is at most $(t/\log t)(\log \log t)^2 \cdot (3/t) \leq (1/\log n)(15k \log \log n)^2$. \square

By Corollary 2.3, even after removing a fraction of $1 - 4/k$ of the edges of $G_{n,k}$, the remaining graph has a connected component with edge expansion at most $O((1/\log n)(k \log \log n)^2)$. So for example, taking $k = 8$ gives a bounded-degree graph such that even after removing up to half of its edges, the remaining graph has a connected component with edge expansion at most $O((1/\log n)(\log \log n)^2)$. At the other extreme, if $k = \log \log n / 648$ then even after removing a $(1 - o(1))$ -fraction of the edges, the remaining graph has a connected component with edge expansion at most $(1/\log n)(\log \log n)^4$.

2.3 Threshold-rank decomposition

Recall that in Section 2.2 we discussed a paradigm for solving unique games by decomposing any graph into expanding subgraphs. The same paradigm was used in the breakthrough paper of Arora, Barak and Steurer [3] which gave the first algorithm to come close to refuting the Unique Games Conjecture, solving $\text{UG}_q(\delta)$ in subexponential time $\exp(q \cdot n^{\delta^{1/3}})$ (here n is the number of variables). The algorithm in [3] employed the *eigenspace enumeration method* where, unlike Trevisan's

⁵We note that prior to [31], Goldreich and Ron [10] have implicitly proved a result of the same spirit, whose exact quantitative properties are somewhat more complicated to state. See also [13]

⁶Trevisan originally used the notion of conductance; our lower bound applies even for edge expansion.

decomposition, the pseudorandom property which the connected components must satisfy depends on the entire spectrum of their adjacency matrix. As in Section 2.2, we prove that one cannot substantially improve the obtained decomposition. Formal definitions follow.

For a d -regular graph G , denote by $\text{rank}_\tau(G)$ the number (with multiplicities) of eigenvalues λ of the adjacency matrix of G satisfying $|\lambda| > \tau d$. Let $R_n(\eta, \epsilon)$ denote the minimum integer such that for any n -vertex graph G , one can remove at most an ϵ -fraction of its edges so that each connected component H of the new graph satisfies $\text{rank}_{1-\eta}(H^*) \leq R_n(\eta, \epsilon)$, where H^* is obtained from H by adding self-loops to make it $\Delta(G)$ -regular. The main result of [3] can be summarized as follows.

Theorem 5 ([28], Theorem 5.6). *There are absolute constants $c, c' > 0$ such that:*

i. For every $0 < \eta, \epsilon \leq 1$,

$$R_n(\eta, \epsilon) \leq n^{c(\eta/\epsilon^2)^{1/3}}.$$

Moreover, a decomposition achieving this bound can be found in polynomial time.⁷

ii. $\text{UG}_q(\delta)$ can be solved in time $\exp(q \cdot R_n(O(\delta), c'))$.

The first step in the algorithm of [3] (item (i) in Theorem 5 which, as the authors themselves write, is the main component of the algorithm) is the efficient decomposition of any graph, by removing few of its edges, into vertex-disjoint subgraphs of small threshold rank. This decomposition is applied on the graph underlying the input unique game, and then the problem is solved on each connected component. The second step (item (ii) in Theorem 5) follows by applying, for each connected component, a brute force enumeration over the eigenspaces corresponding to its large eigenvalues. The second step is inspired by works of Kolla and Tulsiani [4] and Kolla [16], who previously used eigenspace enumeration to solve unique games on Cayley graphs. Combining the two items of Theorem 5 shows that the algorithm of [3] runs in time $\exp(q \cdot n^{O(\delta^{1/3})})$.

Note that in order to obtain a faster algorithm using the eigenspace enumeration method, one must improve the bound on $R_n(\eta, \epsilon)$ in Theorem 5. Up to the present work it was conceivable⁸ that for example $R_n(\eta, \epsilon) \leq (\log n)^{\text{poly}(\eta)}$, say for $\epsilon \geq \epsilon_0$. Such a decomposition would have implied a quasipolynomial time algorithm for $\text{UG}_q(\delta)$, practically refuting⁹ the Unique Games Conjecture.

As we prove below, the bound on $R_n(\eta, \epsilon)$ in Theorem 5 is essentially tight. Namely, we show that one can derive from Theorem 2 a nearly polynomial lower bound on $R_n(\eta, \epsilon)$. This essentially eliminates the possibility of disproving the Unique Games Conjecture using the eigenspace enumeration approach of [3]. See Section 6 for more details on this. We note that our lower bound applies even when $\epsilon = 1 - o(1)$, that is, even when one is allowed to remove all but a $o(1)$ -fraction of the graph's edges. Naturally, the bound holds regardless of the running time of the decomposition algorithm.

⁷So for example, for every d -regular n -vertex graph G , one can efficiently remove at most (say) 1% of the edges of G so that for each connected component in the resulting graph, its “regularized” adjacency matrix has at most $n^{O(\eta^{1/3})}$ eigenvalues larger than $(1 - \eta)d$.

⁸The main tool used in [3] in order to prove Theorem 5 is a higher-order Cheeger-type inequality. It was recently shown in [5] that this inequality is essentially best possible, but to the best of our knowledge, this does not rule out the possibility of improving Theorem 5 by other means.

⁹Say, under the so-called *Exponential Time Hypothesis*.

Corollary 2.4. *For every $n \geq n_0$, $0 < \eta \leq 1$, and $\epsilon \leq 1 - 3000/\log \log \log n$, we have*

$$R_n(\eta, \epsilon) \geq n^{\eta/(\log \log n)^3}.$$

For the proof of Corollary 2.4 we will need the “easy” direction of the higher-order Cheeger inequality that first appeared in [18]. For completeness, we give a short and simple proof in Section 5.

Lemma 2.5. *Let $G = (V, E)$ be a d -regular (multi-)¹⁰ graph, and let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{|V|}$ be the eigenvalues of the adjacency matrix of G . Then for every k ,*

$$\frac{d - \lambda_k}{2} \leq \min_{S_1, \dots, S_k} \max_{1 \leq i \leq k} \phi_G(S_i),$$

where the minimum is over all collections of k mutually disjoint, non-empty subsets of V .

Proof of Corollary 2.4. Let $G = G_{n,k}$ be a graph from Theorem 1. We will show that any subgraph of G with average degree at least 4 has a connected component H satisfying

$$\text{rank}_{1-\eta}(H^*) \geq \frac{1}{4} n^{\eta/k(200 \log \log n)^2}. \quad (3)$$

Setting $k = \log \log \log n / 648$, this would mean that removing an ϵ -fraction of the edges of G , even with ϵ as large as $1 - 4/k \geq 1 - 3000/\log \log \log n$, leaves a connected component H that, for every $n \geq n_0$, satisfies $\text{rank}_{1-\eta}(H^*) \geq n^{\eta/(\log \log n)^3}$. This would imply the desired lower bound on $R_n(\eta, \epsilon)$.

Consider a subgraph of G with average degree at least 4 and note that it has a connected component H with average degree at least 4. Writing t for the number of vertices of H , Claim 2.2 implies that $t \geq n^{1/72k^2} (\geq \log n / (6k)^2)$. Apply Theorem 2 on H with $\mu = 1/t^{\eta k / (20 \log \log t)^2}$. Note that, as required by Theorem 2, we may assume $\mu \geq 1/t$ (as one may verify, recalling $\eta \leq 1$) and $\mu \leq 2/3$ (as otherwise $n^{\eta/k(200 \log \log n)^2} \leq 1/\mu \leq 3/2$, so (3) trivially holds). We thus obtain

$$b := \frac{1}{4\mu} = \frac{1}{4} t^{\eta k / (20 \log \log t)^2} \geq \frac{1}{4} n^{\eta/k(200 \log \log n)^2}$$

mutually disjoint and non-empty subsets of H that each have edge expansion in H at most $(\log(1/\mu)/\log t)(14 \log \log t)^2 < \eta k/2$. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_t$ be the eigenvalues of the adjacency matrix of H^* (recall that H^* is obtained by adding to H sufficiently many self-loops so as to make it $\Delta(G)$ -regular). Note that adding self-loops does not alter the expansion of any subset. Therefore, Lemma 2.5 implies that $\Delta(G) - \lambda_b < \eta k$. It follows that for every $1 \leq i \leq b$ we have $\lambda_i \geq \lambda_b > \Delta(G) - \eta k \geq (1 - \eta)\Delta(G)$. This shows that $\text{rank}_{1-\eta}(H^*) \geq b$, as required in (3), and therefore completes the proof. \square

2.4 Finding a single vertex-expanding subgraph

For a graph G , a subset $S \subseteq V(G)$ is said to have *vertex expansion* α if $|N(S)| = \alpha |S|$, where $N(S)$ denotes the set of vertices outside of S that have a neighbor in S . Motivated by certain extremal problems related to graph minors, it was shown in [27] that every graph G contains a subgraph H , such that H has good vertex expansion properties and almost the same average degree as G .

¹⁰By this we mean that we allow self-loops.

Theorem 6 ([27], Lemma 1.2). *Let $0 < \epsilon \leq 2^{-8}$. Every graph G of average degree k contains a t -vertex subgraph H with average degree at least $(1 - \epsilon)k$ such that every subset of $V(H)$ of size μt with $1/t \leq \mu \leq 1/2$ has vertex expansion in H at least $\epsilon \cdot \log(1/\mu)/\log t(2 \log \log t)^2$.*

Theorem 2, together with item (ii) of Theorem 1 and Fact 1.2, yields the following corollary, which shows that the expansion guarantee in Theorem 6 above is essentially best possible.

Corollary 2.6. *Let G be a graph from Theorem 2, and let H be a t -vertex subgraph of G . Then for every $1/t \leq \mu \leq 2/3$, there is a subset of $V(H)$ of size in $[\mu t/3, \mu t]$ whose vertex expansion in H is at most $(\log(1/\mu)/\log t) \cdot (14 \log \log t)^2$.*

2.5 Hyperfinite families of graphs

A graph is said to be (ϵ, q) -hyperfinite if one can remove an ϵ -fraction of its edges and thus decompose it into connected components of size at most q each. A family of graphs is said to be *hyperfinite* if there is a function q such that for every $\epsilon > 0$, every graph in the family is $(\epsilon, q(\epsilon))$ -hyperfinite. Hyperfinite families of graphs have been extensively studied in recent years, mainly because of their role in the theory of graph limits of sparse graphs (see [24]). Motivated by certain questions related to the design of property-testing algorithms, it was shown in [7] that a hereditary family of graphs in which every graph has a small edge separator must be hyperfinite. More precisely, the following holds.

Theorem 7 ([7], Corollary 3.2). *Let \mathcal{F} be a hereditary family of graphs such that every n -vertex graph $G \in \mathcal{F}$ satisfies $\text{sep}_E(G) \leq n/\log n(\log \log n)^2$. Then, \mathcal{F} is hyperfinite.*

Using Theorem 1 we show that the edge-separation requirement in Theorem 7 cannot be improved much beyond $n/\log n$.

Corollary 2.7. *The graphs $G_{n,8}$ from Theorem 1 are not $(\frac{1}{2}, n^{1/4608})$ -hyperfinite. In particular, setting*

$$\mathcal{F}_8 = \{G : G \text{ is an induced subgraph of } G_{n,8} \text{ for some } n\}$$

we get a hereditary family of graphs that is not hyperfinite, despite the fact that every n -vertex graph $G \in \mathcal{F}_8$ satisfies $\text{sep}_E(G) \leq (n/\log n)(\log \log n)^2 + 48$.

Proof. To see that the first assertion holds, note that after removing at most half of the edges of $G_{n,8}$, we obtain a graph with average degree at least 4. This graph has a connected component H of average degree at least 4. By Claim 2.2, the number of vertices in H is least $n^{1/4608}$. As to the second assertion of the corollary, note that the first assertion clearly means that \mathcal{F}_8 is not hyperfinite. Also, note that the fact that $\text{sep}_E(G) \leq (n/\log n)(\log \log n)^2 + 48$ holds for every $G \in \mathcal{F}_8$ follows from items (i), (ii) and (iii) of Theorem 1 together with Fact 1.2. \square

3 Proof of Theorem 1

As mentioned earlier, in order to construct the graphs $G_{n,k}$ of Theorem 1 we pick a suitable random subgraph of the hypercube. The d -cube, denoted Q_d , is the graph with vertex set $\{0, 1\}^d$ where two vertices are adjacent if their corresponding vectors differ in exactly one coordinate. Note that Q_d is

a d -regular graph on 2^d vertices. The proof will have three main steps, appearing in the following three subsections. In the first subsection we define the graphs $G_{n,k}$ and prove that they satisfy the first two properties of Theorem 1. In the second step we prove that large subgraphs of the cube do not have good expansion properties. Finally, in the last subsection we combine the results from the first two subsections to prove that the graphs $G_{n,k}$ also satisfy the third property in Theorem 1, thus completing its proof.

3.1 Construction

Our goal in this subsection is to prove the following lemma.

Lemma 3.1. *For any $d \in \mathbb{N}$ and $2 \leq k \leq d$ there is a subgraph $Q = Q_{d,k}$ of the d -cube, on all 2^d vertices, such that:*

- i. The average degree of Q is at least k .*
- ii. The maximum degree of Q is at most $3k$.*
- iii. The girth of Q is at least $d/(3k)^2$.*

For the proof of the above lemma we will need the following upper bound on the number of cycles of a given length in the cube.

Claim 3.2. *The number of cycles of length 2ℓ in the d -cube is at most $2^d(d\ell)^\ell$.*

Proof. We will show that the number of closed walks of length 2ℓ , starting and ending at a given vertex, is at most $(d\ell)^\ell$; this would immediately imply the stated bound. We claim that each such closed walk corresponds to a sequence $(x_1, \dots, x_{2\ell}) \in [d]^{2\ell}$ with the property that $|\{i \in [2\ell] : x_i = t\}|$ is even for every $t \in [d]$. To see this, recall that each vertex in the graph corresponds to an element of $\{0, 1\}^d$. The claim follows by considering, for each edge along a closed walk beginning (and ending) with a given vertex, the unique index in which the bit “flips”. So all that is left is to bound from above the number of sequences as above. Now, observe that all these sequences can be generated by first partitioning the set of indices $\{1, \dots, 2\ell\}$ into ℓ pairs and then assigning to each pair a value from $\{1, \dots, d\}$ (of course, this process will generate some sequences several times). Since the number of ways one can pair the elements of $\{1, \dots, 2\ell\}$ is given by

$$(2\ell - 1)!! = (2\ell - 1)(2\ell - 3) \cdots 3 \cdot 1 \leq 2^\ell \ell! / 2 \leq 2^\ell (\ell/2)^\ell = \ell^\ell ,$$

we get that the number of sequences is bounded from above by $d^\ell \cdot \ell^\ell$. □

To prove the existence of a graph as in Lemma 3.1 we use the so-called probabilistic “deletion method”.

Proof of Lemma 3.1. Denote by G' the random subgraph of the d -cube where each edge is independently retained with probability $p = 3k/d$. Let the random variable X count the number of edges of G' . Then

$$\mathbb{E}[X] = p \cdot d2^{d-1} = 3k \cdot 2^{d-1} . \tag{4}$$

Set $L = d/9k^2$ and let the random variable Y count the number of cycles of length at most L in G' . By Claim 3.2,

$$\mathbb{E}[Y] \leq \sum_{\ell=2}^{L/2} p^{2\ell} 2^d (d\ell)^\ell = 2^d \sum_{\ell=2}^{L/2} (p^2 d\ell)^\ell = 2^d \sum_{\ell=2}^{L/2} (\ell/L)^\ell \leq 2^d \sum_{\ell=2}^{L/2} (1/2)^\ell \leq 2^{d-1}. \quad (5)$$

Let the random variable Z count the total number of “excess” edges in G' , that is,

$$Z = \sum_{v: \deg(v) > 3k} (\deg(v) - 3k).$$

For each vertex v , the random variable $\deg(v)$ follows the binomial distribution $B(d, 3k/d)$, so $\mathbb{E}[\deg(v)] = 3k$. Note that $\sum_v |\deg(v) - 3k| = \sum_v (3k - \deg(v)) + 2Z$, which means $\mathbb{E}[Z] = \frac{1}{2} \sum_v \mathbb{E}|\deg(v) - 3k|$. By J ensen’s inequality,

$$(\mathbb{E}|\deg(v) - 3k|)^2 \leq \mathbb{E}[(\deg(v) - 3k)^2] = \text{Var}[\deg(v)] \leq 3k,$$

implying that

$$\mathbb{E}[Z] = \frac{1}{2} \sum_v \mathbb{E}|\deg(v) - 3k| \leq \sqrt{3k} \cdot 2^{d-1}. \quad (6)$$

Combining (4), (5), and (6) we get

$$\mathbb{E}[X - Y - Z] \geq (3k - 1 - \sqrt{3k}) \cdot 2^{d-1} \geq k \cdot 2^{d-1}, \quad (7)$$

where the last inequality can be easily checked to hold for any $k \geq 2$. Let Q be obtained from G' by removing an arbitrary edge from each cycle of length at most $d/9k^2$, as well as removing, for each vertex v with $\deg(v) > 3k$, arbitrary $\deg(v) - 3k$ adjacent edges. Clearly, Q satisfies the last two requirements in the statement. Moreover, we have from (7) that the expected average degree of Q is at least k . The existence of a subgraph of the d -cube as required immediately follows. \square

3.2 Expansion in the hypercube

In this subsection we prove that large subgraphs of the d -cube are not good edge expanders. Recall that a graph $G = (V, E)$ is said to have edge expansion α if every subset $S \subseteq V$ with $|S| \leq |V|/2$ satisfies $|\partial(S)| \geq \alpha |S|$ (the edge boundary $\partial(S)$ is the set of edges with exactly one vertex in S). Our goal in this subsection is to prove the following lemma.

Lemma 3.3. *Any (not necessarily induced) subgraph of the d -cube with t vertices and average degree r has edge expansion at most $2r \log d / \log(t/2)$.*

We first need the following easy claim, in which the *distance* between two vertices u, v in a graph G , denoted $\delta_G(u, v)$, is the length of a shortest path connecting u and v .

Claim 3.4. *For every graph G of maximum degree $\Delta > 1$, and every t -vertex subset $S \subseteq V(G)$, the average distance $\sum_{\{u, v\} \in \binom{S}{2}} \delta_G(u, v) / \binom{t}{2}$ is at least $\log(t/2) / (2 \log \Delta)$.*

Proof. Let $v \in S$. We claim that there are at least $t/2$ vertices in S of distance at least $\ell = \log(t/2)/\log \Delta$ from v . Indeed, the number of vertices of distance at most $\ell - 1$ from v is at most $\sum_{i=0}^{\ell-1} \Delta^i < \Delta^\ell \leq t/2$. It follows that the average distance is at least $\ell/2$. \square

We are now ready to prove the main result of this subsection.

Proof of Lemma 3.3. Let $H = (S, E)$ be a t -vertex subgraph of the d -cube Q_d of average degree r , and let α denote the edge expansion of H . We need to show that $\alpha \leq 2r \log d / \log(t/2)$. Note that we may assume $\alpha > 0$ (i.e., that H is connected), as otherwise there is nothing to prove. For each vertex $v \in S$, write $(v_1, \dots, v_d) \in \{0, 1\}^d$ for the corresponding binary vector. Notice $\delta_{Q_d}(u, v) = \sum_{i=1}^d |u_i - v_i|$. Observe that

$$\frac{\sum_{\{u,v\} \in E} \delta_{Q_d}(u, v)}{\sum_{\{u,v\} \in \binom{S}{2}} \delta_{Q_d}(u, v)} = \frac{\sum_{i=1}^d \sum_{\{u,v\} \in E} |u_i - v_i|}{\sum_{i=1}^d \sum_{\{u,v\} \in \binom{S}{2}} |u_i - v_i|} =: \frac{\sum_{i=1}^d x_i}{\sum_{i=1}^d y_i} \geq \min_{i: y_i \neq 0} \frac{x_i}{y_i}, \quad (8)$$

where the inequality holds¹¹ for all non-negative reals $x_1, y_1, \dots, x_d, y_d$. Let $i \in [d]$ achieve the minimum in (8) (which is well defined since $y_i \neq 0$ for some i), and set $T = \{v \in S : v_i = 1\}$. Note that since $y_i > 0$ we have $0 < |T| < |S|$ and that we can in fact assume that $0 < |T| \leq |S|/2$ (otherwise we replace T with $S \setminus T$). Noting that in (8) we have $x_i = |\partial T|$ and $y_i = |T| |S \setminus T|$, we deduce that

$$\frac{\sum_{\{u,v\} \in E} \delta_{Q_d}(u, v)}{\sum_{\{u,v\} \in \binom{S}{2}} \delta_{Q_d}(u, v)} \geq \frac{|\partial T|}{|T| |S \setminus T|} \geq \min_{\substack{S' \subseteq S: \\ 0 < |S'| \leq |S|/2}} \frac{|\partial S'|}{|S'| |S \setminus S'|} \geq \frac{\alpha}{|S| - 1}.$$

Therefore,

$$\frac{1}{\binom{|S|}{2}} \sum_{\{u,v\} \in \binom{S}{2}} \delta_{Q_d}(u, v) \leq \frac{2}{\alpha |S|} \sum_{\{u,v\} \in E} \delta_{Q_d}(u, v) = \frac{2|E|}{\alpha |S|} = \frac{r}{\alpha}. \quad (9)$$

Applying Claim 3.4 with $G = Q_d$ and the set S , the left hand side of (9) is at least $\log(t/2)/2 \log d$. The desired bound on α follows. \square

3.3 Putting it all together

To prove Theorem 1 we will show that the graphs constructed in Lemma 3.1 have the property that every subgraph H has a small edge separator. To this end, we start by proving that every subgraph H has a small edge expansion. In the proof we will consider two cases, depending on the number of vertices of H . If H has at most $2^{d^{1/3}}$ vertices, then having high expansion implies (see Lemma 3.5 below) the existence of a cycle of length $O(d/k^2)$; however, such a short cycle does not exist in the graph. Otherwise, since the maximum degree of the graph is bounded, Lemma 3.3 implies that H has small edge expansion. Finally, we boost this “hereditary” non-expansion property to construct a small separator in H (see Lemma 3.6 below).

To execute the above proof strategy we will need the following two lemmas, whose proofs appear in Section 5. Recall that we use $\text{sepe}_E(G)$ to denote the size of the smallest edge separator in G .

¹¹Indeed, denoting the minimum by m , we have $x_i \geq m y_i$ for every i (even if $y_i = 0$), so $(\sum_{i=1}^d x_i) \geq m(\sum_{i=1}^d y_i)$.

Lemma 3.5. *Every connected n -vertex graph with edge expansion α (> 0) that is not a tree has girth at most $2(2/\alpha + 1)(\ln(n) + 1) + 1$. In particular, if $\alpha \leq 1$ then the girth is at most $12 \ln n / \alpha$.*

Lemma 3.6. *Let G be an n -vertex graph whose every t -vertex subgraph has expansion at most $f(t)$, for every $t \geq n/3$ and where $f : [n/3, n] \rightarrow \mathbb{R}$ is decreasing. Then $\text{sep}_E(G) \leq (2/3)nf(n/3)$.*

We can finally give a proof for the main theorem.

Proof of Theorem 1. First, suppose that n is a power of 2, and write $n = 2^d$. Let $G = Q_{d,2k}$ be a subgraph of the d -cube as guaranteed by Lemma 3.1. For the rest of the proof set $t_0 = d/(6k)^2$. The assertion of Lemma 3.1 implies that G has average degree at least $2k$, maximum degree at most $6k$, and girth at least t_0 . So to complete the proof (for $n = 2^d$) we only need to establish item (iii) of Theorem 1. We will start by proving that every t -vertex subgraph of G on at least t_0 vertices has edge expansion at most $\frac{1}{3}(1/\log t)(\log \log t)^2$. We will then “boost” this fact using Lemma 3.6 and get that every t -vertex subgraph as above in fact has an edge separator of size at most $(t/\log t)(\log \log t)^2$, as required. Note that the statement of the theorem implies that $d \geq 2^{2 \cdot 648}$. One consequence of this fact is that $d \geq (\log d)^6$. We will also use the fact that the theorem’s assumption that $k \leq \log d/648$ together with the previous observation implies that $d^{1/3}/(6k)^2 \geq 12$ and that $18k \leq t_0/(3 \log t_0)$. Finally, suppose H' is a subgraph of G and that H' is a forest. As G has maximum degree at most $6k$, we have by Fact 1.2 and the above relation between k and t_0 that $\text{sep}_E(H') \leq \Delta(G) \leq 6k \leq \frac{1}{9}t_0/\log t_0$. Hence, if H' has $t \geq t_0$ vertices then its edge expansion is at most $(\frac{1}{9}t/\log t)/(\frac{1}{3}t) = \frac{1}{3}(1/\log t)$.

Let then H be a t -vertex subgraph of G with $t \geq t_0$ vertices, and let us prove that the edge expansion of H is at most $\frac{1}{3}(1/\log t)(\log \log t)^2$. Since we have already established this fact when H is a forest at the end of the previous paragraph, we assume for the rest of this paragraph that H is not a forest. Suppose first that $t < 2^{d^{1/3}}$. By Lemma 3.5, if the edge expansion of H is at least $1/\log t$ then its girth is at most $12(\log t)^2 < 12d^{2/3} \leq d/(6k)^2$. However, G does not contain a cycle this short. Therefore, the edge expansion of H is indeed at most $\frac{1}{3}(1/\log t)(\log \log t)^2$. Suppose now that $t \geq 2^{d^{1/3}}$. By Lemma 3.3 and the fact that the average degree of H is at most $6k$, the edge expansion of H is at most $24k \log d / \log t$. Thus, to prove our claim it suffices to show that $\sqrt{72k \log d} \leq \log \log t$. This indeed follows from the theorem’s assumption that $k \leq \log d/648$, since it implies that

$$\sqrt{72k \log d} \leq (\log d)/3 = \log(d^{1/3}) \leq \log \log t.$$

Having established that every subgraph of G on at least t_0 vertices has small edge expansion, we now wish to show that every such graph has a small separator. So let H be an arbitrary t -vertex subgraph of G with $t \geq t_0$. We apply Lemma 3.6 on the graph H with the function $f(x) = \frac{1}{3}(1/\log x)(\log \log x)^2$, noting that, as required by the lemma, every t' -vertex subgraph of H with $t' \geq n/3$ ($\geq t_0$) has expansion at most $f(t')$, and $f : [n/3, n] \rightarrow \mathbb{R}$ is decreasing (here we use the fact that, as can be easily checked, $f(t)$ is decreasing for $t \geq 256$, and that $n/3 \geq 256$). We deduce that $\text{sep}_E(H) \leq (2/3)tf(t/3) \leq (t/\log t)(\log \log t)^2$, as required by item (iii) in Theorem 1. This completes the proof of the theorem for the case $n = 2^d$.

Finally, let us consider the case of an arbitrary n , that is, not necessarily a power of 2. In this case we set $d = \lceil \log n \rceil$. Since $n > 2^d/2$, a random n -vertex subgraph of $Q_{d,2k}$ will have, with positive probability, average degree at least k . Let G be such a graph. Then the maximum degree

of G is still at most $6k$, its girth is at least $\log n/(6k)^2$, and it is easy to see that every t -vertex subgraph H of G with $t \geq \log n/(6k)^2$ satisfies $\text{sep}_E(H) \leq (t/\log t)(\log \log t)^2$. This completes the proof. \square

4 Proof of Theorem 2

In the first part of the proof we will show that for every t -vertex subgraph H of G with $t \geq \frac{1}{2} \log n/(6k)^2$, and every $1/t \leq \mu \leq 2/3$, there is a subset $A \subseteq V(H)$ of size $\mu t/3 \leq |A| \leq \mu t$ satisfying

$$\phi_H(A) \leq \frac{\log(1/\mu)}{\log t} \cdot (4 \log \log t)^2. \quad (10)$$

In the second part of the proof we will show that if $t \geq \log n/(6k)^2$ (as guaranteed by the statement) and $1/t \leq \mu \leq 2/3$ then there are in fact $1/4\mu$ mutually disjoint subsets of the same size and edge expansion only slightly larger than in (10).

Before we begin, we collect some useful facts. One can easily verify that the theorem's assumption on k implies that $t \geq \sqrt{\log n}$ ($\geq 2^{16}$),¹² and therefore $\Delta(H) \leq \frac{1}{2} \log \log \log n \leq \log \log t$. Notice we may assume $\mu \geq 1/\sqrt{t}$, as otherwise (10) trivially holds by taking A to be any single vertex, in which case $\phi(A) \leq \Delta(H) \leq \log \log t \leq (\log(1/\mu)/\log t)(4 \log \log t)^2$, as needed. For the rest of the proof we set

$$f(x) = (1/\log x) \cdot (\log \log x)^2.$$

One can check that $f(x)$ is decreasing for $x \geq 256$. Throughout the proof, unless otherwise mentioned, we write $\phi(\cdot) = \phi_H(\cdot)$ for the edge expansion in H , and $\partial(\cdot) = \partial_H(\cdot)$ for the edge boundary in H .

In order to obtain a subset $A \subseteq V(H)$ satisfying (10) we will iteratively find smaller and smaller subsets, such that in each step the edge expansion (in H) of our subset does not increase by much. Formally, we claim that for any set $S \subseteq V(H)$ with $|S| \geq \sqrt{t}$ (≥ 256) there is a subset $S_1 \subseteq S$ of size $\frac{1}{3}|S| \leq |S_1| \leq \frac{2}{3}|S|$ such that $\phi(S_1) \leq \phi(S) + 2f(|S|)$. To see this, consider an edge separator in the induced subgraph $H[S]$. Namely, let $S = S' \cup S''$ be a partition of S with $|S'|, |S''| \leq \frac{2}{3}|S|$ such that the number of edges between S' and S'' in $H[S]$, and thus in H , is $\text{sep}_E(H[S])$. Consider now the edge boundary of S' and S'' in H . We have

$$|\partial(S')| + |\partial(S'')| = |\partial(S)| + 2\text{sep}_E(H[S]).$$

Moving to edge expansion, observe that this means that

$$\min\{\phi(S'), \phi(S'')\} \leq \frac{|\partial(S')| + |\partial(S'')|}{|S'| + |S''|} = \frac{|\partial(S)| + 2\text{sep}_E(H[S])}{|S|} \leq \phi(S) + 2f(|S|),$$

where in the left inequality we used the elementary inequality $(a+b)/(c+d) \geq \min\{a/c, b/d\}$, and in the right inequality we used the fact that $\text{sep}_E(H[S]) \leq |S|f(|S|)$. This last inequality follows from item (iii) of Theorem 1 when $|S| \geq \log n/(6k)^2$, while when $|S| < \log n/(6k)^2$ it follows from item (ii) of Theorem 1, Fact 1.2 and the fact that f is decreasing, as these facts together mean

¹² $72k^2 \leq (\log \log \log n)^2 \leq \sqrt{\log n}$; for the last inequality, note the statement of the theorem implies $\log n \geq 2^{2 \cdot 648}$.

that $\text{sep}_E(H[S]) \leq \Delta(H) \leq \log \log t \leq \sqrt{t}f(t) \leq |S|f(|S|)$. Altogether we see that we can take either S' or S'' as the set S_1 . Applying this claim with $S = S_0 = V(H)$ yields a subset $S_1 \subseteq S_0$ of size $\frac{1}{3}|S_0| \leq |S_1| \leq \frac{2}{3}|S_0|$ satisfying $\phi(S_1) \leq \phi(S_0) + 2f(|S_0|) = 2f(|S_0|)$, where here we use the fact that $S_0 = V(H)$ so $\phi(S_0) = 0$. Having picked S_i , for some $i \geq 1$, we can then pick a new set $S_{i+1} \subseteq S_i$ satisfying $\frac{1}{3}|S_i| \leq |S_{i+1}| \leq \frac{2}{3}|S_i|$ and $\phi(S_{i+1}) \leq \phi(S_i) + 2f(|S_i|)$. We stop the process once we obtain a subset $A \subseteq V(H)$ of size $|A| \leq \mu t$. Note that this means that $|A| \geq \mu t/3$. Let $V(H) = S_0 \supseteq S_1 \supseteq \dots \supseteq S_s = A$ denote the subsets we thus obtain. We get from the above observations that

$$\phi(A) \leq \sum_{i=0}^{s-1} 2f(|S_i|) \leq s \cdot 2f(\mu t),$$

where in the second inequality we used the fact that $f(x)$ is decreasing. Since the number s of iterations is at most $\lceil \log_{3/2}(1/\mu) \rceil \leq 2 \log_{3/2}(1/\mu)$, we get

$$\phi(A) \leq 2 \log_{3/2}(1/\mu) \cdot 2f(\sqrt{t}) \leq \frac{\log(1/\mu)}{\log t} \cdot 16(\log \log t)^2,$$

where in the first inequality we used the assumption that $\mu t \geq \sqrt{t}$. This completes the first part of the proof.

Having found a single non-expanding subset A satisfying (10), we proceed to find at least $1/4\mu$ disjoint subsets of size in $[\mu t/3, \mu t]$ having small expansion, as required by the theorem's statement. So let $H = (V, E)$ be a t -vertex subgraph with $t \geq \log n/(6k)^2$ and suppose $1/t \leq \mu \leq 2/3$. Since $|V| \geq \frac{1}{2} \log n/(6k)^2$, we can find a subset $A_1 \subseteq V$ satisfying $\phi_H(A_1) \leq \varphi(\mu, t)$ where $\varphi(\mu, t)$ is the right hand side of (10). Now apply the same argument on the induced subgraph $H[V \setminus A_1]$ in order to find a subset $A_2 \subseteq V \setminus A_1$ of size in $[\mu t/3, \mu t]$ that satisfies¹³ $\phi_{H[V \setminus A_1]}(A_2) \leq \varphi(\mu, t)$. Note that we can pick sets $A_i \subseteq V \setminus \bigcup_{j=1}^{i-1} A_j$ at least $1/2\mu$ times (assume henceforth that $1/2\mu$ is an integer), since as long as we pick fewer sets we are left with at least $t/2 \geq \frac{1}{2} \log n/(6k)^2$ vertices, so we can still pick another set. We therefore obtain at least $1/2\mu$ mutually disjoint subsets $A_1, \dots, A_{1/2\mu} \subseteq V$ of size in $[\mu t/3, \mu t]$ each, satisfying

$$\phi_{H[V \setminus \bigcup_{j=1}^{i-1} A_j]}(A_i) \leq \varphi(\mu, t) \quad (11)$$

for every $1 \leq i \leq 1/2\mu$. To finish the proof, we claim that at least half of the $1/2\mu$ subsets A_i have small expansion in H itself. To see this, note that

$$\sum_{i=1}^{1/2\mu} |\partial_H(A_i)| \leq 2 \sum_{i=1}^{1/2\mu} |\partial_{H[V \setminus \bigcup_{j=1}^{i-1} A_j]}(A_i)| = 2 \sum_{i=1}^{1/2\mu} |A_i| \cdot \phi_{H[V \setminus \bigcup_{j=1}^{i-1} A_j]}(A_i) \leq t\varphi(\mu, t), \quad (12)$$

where in the first inequality we use the fact that each edge is counted (at most) twice, and in the second inequality we use (11) and the fact that each A_i is of size at most μt . Dividing both sides of (12) by $\mu t/3$ (which is an upper bound on the size of the sets A_i) we get that $\sum_{i=1}^{1/2\mu} \phi_H(A_i) \leq 3\varphi(\mu, t)/\mu$. Hence, by averaging, at least $1/4\mu$ of the sets A_i satisfy

$$\phi_H(A_i) \leq 12\varphi(\mu, t) \leq (\log(1/\mu)/\log t)(14 \log \log t)^2,$$

thus completing the proof.

¹³A subset $A_2 \subseteq V \setminus A_1$ of size at most μt is therefore of size at most $\mu' |V \setminus A_1|$ with $\mu' \geq \mu$. Thus, we can find a subset A_2 satisfying $\phi_{H[V \setminus A_1]}(A_2) \leq \varphi(\mu', t) \leq \varphi(\mu, t)$ where we use the monotonicity of $\varphi(\mu, t)$ with respect to μ .

5 Missing Proofs

5.1 Proof of Lemma 3.5

Proof of Lemma 3.5. Let k denote the average degree of our graph $G = (V, E)$. Since G is connected and not a tree we have $k \geq 2$. As is well known, G must have a subgraph of minimum degree at least $\lfloor k/2 \rfloor + 1 \geq 2$. In particular, this subgraph contains a cycle. Let $e = uv$ be any edge in such a cycle, and let G' be obtained from G by removing e . Note that our choice of e guarantees that $\deg_{G'}(u), \deg_{G'}(v) \geq \lfloor k/2 \rfloor$. We claim that the edge expansion of G' is at least $\alpha/2$. Indeed, let $S \subseteq V$ with $e \in \partial_G(S)$. Then $|\partial_G(S)| \geq 2$, as otherwise S forms a connected component of G' , contradicting the fact that e is not a cut edge in G . Therefore, removing a single edge can decrease the edge expansion of G by at most a factor of 2.

For a vertex $x \in V$ and a non-negative integer r , write $B_x(r) = \{y : \delta_{G'}(x, y) \leq r\}$. Let $m_x(r)$ denote the number of edges spanned by $B_x(r)$. Since G' has edge expansion at least $\alpha/2$, we have $m_x(r+1) \geq (1 + \alpha/2)m_x(r)$ provided $|B_x(r)| \leq n/2$. Put $R = (2/\alpha + 1)(\ln(n) + 1)$. We claim that $|B_v(R)| > n/2$. Indeed, recalling $m_v(1) = \deg_{G'}(v) \geq \lfloor k/2 \rfloor$, we get that if $|B_v(R)| \leq n/2$ then

$$m_v(R+1) \geq (1 + \alpha/2)^R \cdot \lfloor k/2 \rfloor \geq e^{\ln(n)+1} \cdot \lfloor k/2 \rfloor \geq nk/2,$$

which is strictly greater than the number of edges of G' – a contradiction. Since we also have $\deg_{G'}(u) \geq \lfloor k/2 \rfloor$ we deduce that $B_u(R) \cap B_v(R) \neq \emptyset$. This means that $\delta_{G'}(u, v) \leq 2R$. Since u and v are adjacent in G but not in G' , we conclude that G has a cycle of length at most $2R + 1 = 2(2/\alpha + 1)(\ln(n) + 1) + 1$. \square

5.2 Proof of Lemma 3.6

Proof of Lemma 3.6. We iteratively remove subsets $S_1, \dots, S_k \subseteq V(G)$ as follows. For the i -th step, consider the graph $G_i := G - \bigcup_{j=1}^{i-1} S_j$ and let $S_i \subseteq V(G_i)$ have size $|S_i| \leq |V(G_i)|/2$ and expansion in G_i at most $f(|V(G_i)|)$. We stop removing subsets once the set $S := \bigcup_{j=1}^k S_j$ is of size at least $n/3$, and thus $|S| \leq \sum_{j=1}^{k-1} |S_j| + |V(G_k)|/2 = (n + \sum_{j=1}^{k-1} |S_j|)/2 \leq 2n/3$. Now, observe that every edge in the edge boundary $\partial_G(S)$ is a member of some edge boundary $\partial_{G_j}(S_j)$. Hence, the expansion of S in G is

$$\frac{|\partial_G(S)|}{|S|} \leq \frac{\sum_{j=1}^k |\partial_{G_j}(S_j)|}{\sum_{j=1}^k |S_j|} \leq \max_{1 \leq j \leq k} \frac{|\partial_{G_j}(S_j)|}{|S_j|} \leq \max_{1 \leq j \leq k} f(|V(G_j)|) \leq f(n/3),$$

where in the second inequality we used the inequality $(\sum_{j=1}^k x_j)/(\sum_{j=1}^k y_j) \leq \max_{1 \leq j \leq k} (x_j/y_j)$ for non-negative reals x_j, y_j ,¹⁴ and in the last inequality we used the fact that $|V(G_j)| \geq |V(G_k)| \geq n - |S| \geq n/3$ together with the fact that $f(t)$ is decreasing for $t \geq n/3$. It follows that $|\partial_G(S)| \leq |S| f(n/3) \leq (2/3)nf(n/3)$, completing the proof. \square

¹⁴We note that while the simple argument here suffices to bound the expansion of $\bigcup S_i$, bounding the expansion of each individual S_i – which is done in the second part of the proof of Theorem 2 – is trickier.

5.3 Proof of Lemma 2.5

Proof of Lemma 2.5. Put $L = dI - A$ where A is the adjacency matrix of $G = (V, E)$. Note that the k -th eigenvalue of L is $d - \lambda_k$. By the Courant-Fischer well-known min-max theorem, we have

$$d - \lambda_k = \min_{w^1, \dots, w^k} \max_{0 \neq x \in \text{span}\{w^1, \dots, w^k\}} R_L(x), \quad (13)$$

where the minimum is over all collections of k mutually orthogonal, nonzero vectors in $\mathbb{R}^{|V|}$, and $R_L(x) = x^t L x / x^t x$ is the Rayleigh quotient. Let $S_1, \dots, S_k \subseteq V$ be mutually disjoint, non-empty subsets, and denote by $\mathbf{1}_{S_i}$ the characteristic $\{0, 1\}$ -vector of S_i . As it is easy to see that $x^t L y = \sum_{\{u, v\} \in E} (x_u - x_v)(y_u - y_v)$, we deduce that $\mathbf{1}_{S_i}^t L \mathbf{1}_{S_i} = |\partial(S_i)|$, and that for $i \neq j$ we have $\mathbf{1}_{S_i}^t L \mathbf{1}_{S_j} = -e(S_i, S_j)$ where $e(S_i, S_j)$ is the number of edges between S_i and S_j . Consider now any vector $x \in \text{span}\{\mathbf{1}_{S_1}, \dots, \mathbf{1}_{S_k}\}$ and write $x = \sum_{i=1}^k c_i \mathbf{1}_{S_i}$. Then

$$\begin{aligned} x^t L x &= \sum_{i,j=1}^k c_i c_j \cdot \mathbf{1}_{S_i}^t L \mathbf{1}_{S_j} = \sum_{i=1}^k c_i^2 |\partial(S_i)| - \sum_{i \neq j} c_i c_j e(S_i, S_j) \leq \sum_{i=1}^k c_i^2 |\partial(S_i)| + \sum_{i \neq j} \frac{c_i^2 + c_j^2}{2} e(S_i, S_j) \\ &= \sum_{i=1}^k c_i^2 |\partial(S_i)| + \sum_{i=1}^k c_i^2 \sum_{j \neq i} e(S_i, S_j) \leq 2 \sum_{i=1}^k c_i^2 |\partial(S_i)|. \end{aligned}$$

The proof now follows since (13) implies

$$d - \lambda_k \leq \max_{0 \neq x \in \text{span}\{\mathbf{1}_{S_1}, \dots, \mathbf{1}_{S_k}\}} R_L(x) \leq \frac{2 \sum_{i=1}^k c_i^2 |\partial(S_i)|}{\sum_{i=1}^k c_i^2 |S_i|} \leq 2 \max_{1 \leq i \leq k} \phi_G(S_i).$$

□

6 Concluding Remarks and Open Problems

- As we mentioned in Section 1, we can replace the $(\log \log t)^2$ term in (1) with $(\log \log t)^{1+o(1)}$. It would be interesting to know whether the $\log \log t$ factors in (1) (and thus also in (2)) can be removed. One way to obtain this is to improve Lemma 3.3. In this regard, we conjecture the following “higher-order” isoperimetric property of the hypercube: Among all subgraphs of the d -cube on 2^i vertices, for every $i \in \mathbb{N}$ at least $d^{1/3}$ say, the subcubes have the largest normalized¹⁵ edge expansion. (One can also come up with an analogous isoperimetric conjecture for vertex expanders in the hypercube.) Since a subcube on $t = 2^i$ vertices has normalized edge expansion $1/i = 1/\log t$, proving the above conjecture would mean that the additional $\log d$ factor in Lemma 3.3 is unnecessary for $t \geq 2^{d^{1/3}}$ (note that without the lower bound on t , Lemma 3.3 is tight as witnessed by $K_{1,d}$). By the proof of Theorem 1, this would mean that the graphs $G_{n,k}$ are such that the $\log \log t$ terms in (1) and (2) can in fact be replaced by $O(k)$. Note that this is best possible, as witnessed by the subgraph $G_{n,k}$ itself. We note that this would essentially close the gaps in the applications mentioned in Section 2, as k is chosen there to be either constant or grow arbitrarily small with n . For example,

¹⁵I.e., edge expansion divided by the subgraph’s average degree.

this would imply a truly polynomial lower bound on the threshold rank in Corollary 2.4, thus proving that the eigenspace enumeration approach of [3] cannot solve $\text{UG}_q(\delta)$ in time $\exp(n^{o(1)})$ for constant δ, q .

- It is not hard to see that if every t -vertex subgraph H of a graph G satisfies $\text{sep}_E(H) \leq (t/\log t)(\log \log t)^c$ then G must have at most $O(n(\log \log n)^{c'})$ edges. (Heuristically speaking, the upper bound on the number of edges admits a recursion of the form $f(n) \leq 2f(n/2) + (n/\log n)(\log \log n)^c$). We note that one can construct a graph G satisfying this requirement using arguments similar to those we used to prove Theorem 1, and in this case, the resulting graph would have $\Omega(n(\log \log n)^{c''})$ edges. (To be more precise, $c' = c + 1$ and $c'' = c - 1$.) This means that the graph we construct in Theorem 1 has the maximum possible number of edges, up to $\log \log n$ factors.

The situation for vertex separators, as in Theorem 3, is not so clear, so it would be interesting to understand the maximum number of edges of n -vertex graphs in a hereditary family with vertex separators of size at most $(n/\log n)(\log \log n)^c$. Although we construct such a family \mathcal{F} that contains n -vertex graphs with $\Omega(n(\log \log n)^{c''})$ edges, we cannot rule out the possibility of improving this to (say) $n \log n$.

- Let us mention two problems that seem to be related to the types of problems studied here but that (unfortunately) we cannot resolve. The first is a nice problem of G. Kalai. Suppose \mathcal{F} is a hereditary family of graphs and that every graph in the family has a vertex separator of size $n/f(n)$. Then, how fast should $f(n)$ grow so that \mathcal{F} has only $2^{O(n)}$ non-isomorphic graphs on n vertices? It was shown in [11, 12] that $f(n) \geq \log^{2+\epsilon} n$ suffices while $f(n) \leq \log^{1-\epsilon} n$ does not. Closing the gap between these bounds is still open.

The second problem is related to a graph decomposition result from a work of Pătraşcu and Thorup [25]. They proved that the edge set of every graph can be decomposed into $b = O(\log n)$ subsets E_1, \dots, E_b so that for every $1 \leq i \leq b$ the graph spanned by E_1, \dots, E_i has “edge expansion” $1/\log n$, where the notion of edge expansion used here is slightly different from the (standard) one we used throughout this paper. It would be interesting to decide whether the parameters in the construction of [25] are optimal.

Acknowledgment: We are grateful to Alex Samorodnitsky for very helpful discussions related to this work.

References

- [1] N. Alon, S. Hoory and N. Linial, The Moore bound for irregular graphs, *Graph. Combinator.* **18** (2002), 53–57. 2.2
- [2] N. Alon, P. Seymour and R. Thomas, Planar separators, *SIAM J. Discrete Math.* **7** (1994), 184–193. 2.1
- [3] S. Arora, B. Barak and D. Steurer, Subexponential algorithms for unique games, *Proc. FOCS 2010*, 563–572. 1, 1.1, 2.2, 2.3, 2.3, 8, 6

- [4] S. Arora, A. Kolla, S. Khot, D. Steurer, M. Tulsiani and N. Vishnoi, Unique games on expanding constraint graphs are easy, *Proc. STOC 2008*, 21–28. [2.3](#)
- [5] B. Barak, P. Gopalan, J. Håstad, R. Meka, P. Raghavendra and D. Steurer, Making the long code shorter, *Proc. FOCS 2012*, 370–379. [8](#)
- [6] D. Conlon, J. Fox and B. Sudakov, Cycle packing, Manuscript, 2013. [1](#)
- [7] A. Czumaj, A. Shapira and C. Sohler, Testing hereditary properties of non-expanding bounded-degree graphs, *SIAM J. Comput.* **38** (2009), 2499–2510. [1](#), [2.5](#), [7](#)
- [8] J. Fox and J. Pach, Separator theorems and Turán-type results for planar intersection graphs, *Adv. Math.* **219** (2008), 1070–1080. [2.1](#), [3](#)
- [9] A. Frieze and R. Kannan, Quick approximation to matrices and applications, *Combinatorica* **19** (1999), no. 2, 175–220. [1](#)
- [10] O. Goldreich and D. Ron, A sublinear bipartiteness tester for bound degree graphs, *Combinatorica* **19** (1999), 335–373. [1](#), [5](#)
- [11] M. Litvin, Counting graphs, counting edges and separation properties, M.Sc Thesis, Jerusalem 1996. [6](#)
- [12] G. Kalai, Separator theorems and small hereditary classes (Reserach Problem), *Discrete Mathematics* **257** (2002), 611–612. [6](#)
- [13] R. Kannan, S. Vempala and A. Vetta, On clusterings: Good, bad and spectral, *J. ACM* **51** (2004), 497–515. [5](#)
- [14] S. Khot, On the power of unique 2-prover 1-round games, *Proc. STOC 2002* **34** (2002), 767–775. [1](#), [2.2](#), [4](#)
- [15] S. Khot, G. Kindler, E. Mossel and R. O’Donnell, Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?, *SIAM J. Comput.*, **37** (2007), 319–357. [4](#)
- [16] A. Kolla, Spectral Algorithms for Unique Games, *Proc. CCC 2010*, 122–130. [2.3](#)
- [17] A. Kostochka and L. Pyber, Small topological complete subgraphs of “dense” graphs, *Combinatorica* **8** (1988), 83–86. [1](#)
- [18] J. Lee, S. Oveis Gharan and L. Trevisan, Multi-way spectral partitioning and higher-order Cheeger inequalities, *J. ACM*, to appear. [2.3](#)
- [19] F. Leighton and S. Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, *J. ACM* **46** (1999), 787–832. [1](#)
- [20] N. Linial, E. London and E. Rabinovich, The geometry of graphs and some of its algorithmic applications, *Combinatorica* **15** (1995), 215–245. [1.2](#)

- [21] N. Linial and M. Saks, Low diameter graph decompositions, *Combinatorica* **13** (1993), 441–454. [1](#)
- [22] R. Lipton, D. Rose and R. Tarjan, Generalized nested dissections, *SIAM J. Numer. Anal.* **16** (1979), no. 2, 346–358. [1](#), [2.1](#), [3](#)
- [23] R. Lipton and R. Tarjan, A separator theorem for planar graphs, *SIAM J. Appl. Math* **36** (1979), no. 2, 177–189. [2.1](#)
- [24] L. Lovász, **Large networks and graph limits**, AMS, 2012. [2.5](#)
- [25] M. Pătraşcu and M. Thorup, Planning for fast connectivity updates, *Proc. FOCS 2007*, 263–271. [1](#), [6](#)
- [26] A. L. Rosenberg and L. S. Heath, **Graph Separators, with Applications**, Springer, 2000. [2.1](#)
- [27] A. Shapira and B. Sudakov, Small complete minors above the extremal edge density, *Combinatorica*, to appear. [1](#), [2.4](#), [6](#)
- [28] D. Steurer, On the complexity of unique games and graph expansion, Dissertation, 2010. [1.1](#), [5](#)
- [29] D. Steurer and P. Raghavendra, Graph expansion and the unique games conjecture, *Proc. STOC 2010*, 755–764. [1.1](#)
- [30] E. Szemerédi, Regular partitions of graphs, In: *Proc. Colloque Inter. CNRS* (J. C. Bermond, J. C. Fournier, M. Las Vergnas and D. Sotteau, eds.), 1978, 399–401. [1](#)
- [31] L. Trevisan, Approximation algorithms for unique games, *Proc. FOCS 2005*, 197–205. [1](#), [2.2](#), [4](#), [5](#)